Postel

Network Working Group
Request for Comments: 1017

Barry M. Leiner
RIACS
August 1987

Network Requirements for Scientific Research

Internet Task Force on Scientific Computing

STATUS OF THIS MEMO

This RFC identifies the requirements on communication networks for supporting scientific research. It proposes some specific areas for near term work, as well as some long term goals. This is an "idea" paper and discussion is strongly encouraged. Distribution of this memo is unlimited.

INTRODUCTION

Computer networks are critical to scientific research. They are currently being used by portions of the scientific community to support access to remote resources (such as supercomputers and data at collaborator's sites) and collaborative work through such facilities as electronic mail and shared databases. There is considerable movement in the direction of providing these capabilities to the broad scientific community in a unified manner, as evidence by this workshop. In the future, these capabilities will even be required in space, as the Space Station becomes a reality as a scientific research resource.

The purpose of this paper is to identify the range of requirements for networks that are to support scientific research. These requirements include the basic connectivity provided by the links and switches of the network through the basic network functions to the user services that need to be provided to allow effective use of the interconnected network. The paper has four sections. The first section discusses the functions a user requires of a network. The second section discusses the requirements for the underlying link and node infrastructure while the third proposes a set of specifications to achieve the functions on an end-to-end basis. The fourth section discusses a number of network-oriented user services that are needed in addition to the network itself. In each section, the discussion is broken into two categories. The first addresses near term requirements: those capabilities and functions that are needed today and for which technology is available to perform the function. The second category concerns long term goals: those capabilities for which additional research is needed.

This RFC was produced by the IAB Task force a Scientific Computing,

which is chartered to investigate advanced networking requirements that result from scientific applications. Work reported herein was supported in part by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration (NASA) to the Universities Space Research Association (USRA).

1. NETWORK FUNCTIONS

This section addresses the functions and capabilities that networks and particularly internetworks should be expected to support in the near term future.

Near Term Requirements

There are many functions that are currently available to subsets of the user community. These functions should be made available to the broad scientific community.

User/Resource Connectivity

Undoubtedly the first order of business in networking is to provide interconnectivity of users and the resources they need. The goal in the near term for internetworking should be to extend the connectivity as widely as possible, i.e. to provide ubiquitous connectivity among users and between users and resources. Note that the existence of a network path between sites does not necessarily imply interoperability between communities and or resources using non-compatible protocol suites. However, a minimal set of functions should be provided across the entire user community, independent of the protocol suite being used. These typically include electronic mail at a minimum, file transfer and remote login capabilities must also be provided.

Home Usage

One condition that could enhance current scientific computing would be to extend to the home the same level of network support that the scientist has available in his office environment. As network access becomes increasingly widespread, the extension to the home will allow the user to continue his computing at home without dramatic changes in his work habits, based on limited access.

Charging

The scientific user should not have to worry about the costs of data communications any more than he worries about voice communications (his office telephone), so that data communications becomes an integral and low-cost part of our national infrastructure. This

implies that charges for network services must NOT be volume sensitive and must NOT be charged back to the individual. Either of these conditions forces the user to consider network resources as scarce and therefore requiring his individual attention to conserve them. Such attention to extraneous details not only detracts from the research, but fundamentally impacts the use and benefit that networking is intended to supply. This does not require that networking usage is free. It should be either be low enough cost that the individual does not have to be accountable for "normal" usage or managed in such a manner that the individual does not have to be concerned with it on a daily basis.

Applications

Most applications, in the near term, which must be supported in an internetwork environment are essentially extensions of current ones. Particularly:

Electronic Mail

Electronic mail will increase in value as the extended interconnectivity provided by internetworking provides a much greater reachability of users.

Multimedia Mail

An enhancement to text based mail which includes capabilities such as figures, diagrams, graphs, and digitized voice.

Multimedia Conferencing

Network conferencing is communication among multiple people simultaneously. Conferencing may or may not be done in "real time", that is all participants may not be required to be on-line at the same time. The multimedia supported may include text, voice, video, graphics, and possibly other capabilities.

File Transfer

The ability to transfer data files.

Bulk Transfer

The ability to stream large quantities of data.

Interactive Remote Login

The ability to perform remote terminal connections to hosts.

Remote Job Entry

The ability to submit batch jobs for processing to remote hosts and receive output.

Applications which need support in the near term but are NOT extensions of currently supported applications include:

Remote Instrument Control

This normally presumes to have a human in the "control loop". This condition relaxes the requirements on the (inter)network somewhat as to response times and reliability. Timing would be presumed to be commensurate with human reactions and reliability would not be as stringent as that required for completely automatic control.

Remote Data Acquisition

This supports the collection of experimental data where the experiment is remotely located from the collection center. This requirement can only be satisfied when the bandwidth, reliability, and predictability of network response are sufficient. This cannot be supported in the general sense because of the enormous bandwidth, very high reliability, and/or guaranteed short response time required for many experiments.

These last two requirements are especially crucial when one considers remote experimentation such as will be performed on the Space Station.

Capabilities

The above applications could be best supported on a network with infinite bandwidth, zero delay, and perfect reliability. Unfortunately, even currently feasible approximations to these levels of capabilities can be very expensive. Therefore, it can be expected that compromises will be made for each capability and between them, with different balances struck between different networks. Because of this, the user must be given an opportunity to declare which capability or capabilities is/are of most interest-most likely through a "type-of-service" required declaration. Some examples of possible trade-offs: File Transport Normally requires high reliability primarily and high bandwidth secondarily. Delay is not as important.

Bulk Transport

Some applications such as digitized video might require high bandwidth as the most important capability. Depending on the application, delay would be second, and reliability of lesser importance. Image transfers of scientific data sometimes will invert the latter two requirements.

Interactive Traffic

This normally requires low delay as a primary consideration. Reliability may be secondary depending on the application. Bandwidth would usually be of least importance.

Standards

The use of standards in networking is directed toward interoperability and availability of commercial equipment. However, as stated earlier, full interoperability across the entire scientific community is probably not a reasonable goal for internetworking in the near term because of the protocol mix now present. That is not to say, though, that the use of standards should not be pursued on the path to full user interoperability. Standards, in the context of near term goal support, include:

Media Exchange Standards

Would allow the interchange of equations, graphics, images, and data bases as well as text.

Commercially Available Standards

Plug compatible, commercially available standards will allow a degree of interoperability prior to the widespread availability of the ISO standard protocols.

Long Term Goals

In the future, the internetwork should be transparent communications between users and resources, and provide the additional network services required to make use of that communications. A user should be able to access whatever resources are available just as if the resource is in the office. The same high level of service should exist independent of which network one happens to be on. In fact, one should not even be able to tell that the network is there!

It is also important that people be able to work effectively while at home or when traveling. Wherever one may happen to be, it should be

possible to "plug into" the internetwork and read mail, access files, control remote instruments, and have the same kind of environment one is used to at the office.

Services to locate required facilities and take advantage of them must also be available on the network. These range from the basic "white" and "yellow" pages, providing network locations (addresses) for users and capabilities, through to distributed data bases and computing facilities. Eventually, this conglomeration of computers, workstations, networks, and other computing resources will become one gigantic distributed "world computer" with a very large number of processing nodes all over the world.

2. NETWORK CONNECTIVITY

By network connectivity, we mean the ability to move packets from one point to another.

Note that an implicit assumption in this paper is that packet switched networks are the preferred technology for providing a scientific computer network. This is due to the ability of such networks to share the available link resources to provide interconnection between numerous sites and their ability to effectively handle the "bursty" computer communication requirement.

Note that this need not mean functional interoperability, since the endpoints may be using incompatible protocols. Thus, in this section, we will be addressing the use of shared links and interconnected networks to provide a possible path. In the next section, the exploitation of these paths to achieve functional connectivity will be addressed.

In this section, we discuss the need for providing these network paths to a wide set of users and resources, and the characteristics of those paths. As in other sections, this discussion is broken into two major categories. The first category are those goals which we believe to be achievable with currently available technology and implementations. The second category are those for which further research is required.

Near Term Objectives

Currently, there are a large number of networks serving the scientific community, including Arpanet, MFEnet, SPAN, NASnet, and the NSFnet backbone. While there is some loose correlation between the networks and the disciplines they serve, these networks are organized more based on Federal funding. Furthermore, while there is significant interconnectivity between a number of the networks, there

is considerable room for more sharing of these resources.

In the near term, therefore, there are two major requirement areas; providing for connectivity based on discipline and user community, and providing for the effective use of adequate networking resources.

Discipline Connectivity

Scientists in a particular community/discipline need to have access to many common resources as well as communicate with each other. For example, the quantum physics research community obtains funding from a number of Federal sources, but carries out its research within the context of a scientific discourse. Furthermore, this discourse often overlaps several disciplines. Because networks are generally oriented based on the source of funding, this required connectivity has in the past been inhibited. NSFnet is a major step towards satisfying this requirement, because of its underlying philosophy of acting as an interconnectivity network between supercomputer centers and between state, regional, and therefore campus networks. This move towards a set of networks that are interconnected, at least at the packet transport level, must be continued so that a scientist can obtain connectivity between his/her local computing equipment and the computing and other resources that are needed, independently of the source of funds.

Obviously, actual use of those resources will depend on obtaining access permission from the appropriate controlling organization. For example, use of a supercomputer will require permission and some allocation of computing resources. The lack of network access should not, however, be the limiting factor for resource utilization.

Communication Resource Sharing

The scientific community is always going to suffer from a lack of adequate communication bandwidth and connections. There are requirements (e.g. graphic animation from supercomputers) that stretch the capabilities of even the most advanced long-haul networks. In addition, as more and more scientists require connection into networks, the ability to provide those connections on a network-centric basis will become more and more difficult.

However, the communication links (e.g. leased lines and satellite channels) providing the underlying topology of the various networks span in aggregate a very broad range of the scientific community sites. If, therefore, the networks could share these links in an effective manner, two objectives could be achieved:

The need to add links just to support a particular network

topology change would be decreased, and

New user sites could be connected more readily.

Existing technology (namely the DARPA-developed gateway system based on the Internet Protocol, IP) provides an effective method for accomplishing this sharing. By using IP gateways to connect the various networks, and by arranging for suitable cost-sharing, the underlying connectivity would be greatly expanded and both of the above objectives achieved.

Expansion of Physical Structure

Unfortunately, the mere interconnectivity of the various networks does not increase the bandwidth available. While it may allow for more effective use of that available bandwidth, a sufficient number of links with adequate bandwidth must be provided to avoid network congestion. This problem has already occurred in the Arpanet, where the expansion of the use of the network without a concurrent expansion in the trunking and topology has resulted in congestion and consequent degradation in performance.

Thus, it is necessary to augment the current physical structure (links and switches) both by increasing the bandwidth of the current configuration and by adding additional links and switches where appropriate.

Network Engineering

One of the major deficiencies in the current system of networks is the lack of overall engineering. While each of the various networks generally is well supported, there is woefully little engineering of the overall system. As the networks are interconnected into a larger system, this need will become more severe. Examples of the areas where engineering is needed are:

Topology engineering-deciding where links and switches should be installed or upgraded. If the interconnection of the networks is achieved, this will often involve a decision as to which networks need to be upgraded as well as deciding where in the network those upgrades should take place.

Connection Engineering-when a user site desires to be connected, deciding which node of which network is the best for that site, considering such issues as existing node locations, available bandwidth, and expected traffic patterns to/from that site.

Operations and Maintenance-monitoring the operation of the overall

system and identifying corrective actions when failures occur.

Support of Different Types of Service

Several different end user applications are currently in place, and these put different demands on the underlying structure. For example, interactive remote login requires low delay, while file transfer requires high bandwidth. It is important in the installation of additional links and switches that care be given to providing a mix of link characteristics. For example, high bandwidth satellite channels may be appropriate to support broadcast applications or graphics, while low delay will be required to support interactive applications.

Future Goals

Significant expansion of the underlying transport mechanisms will be required to support future scientific networking. These expansions will be both in size and performance.

Bandwidth

Bandwidth requirements are being driven higher by advances in computer technology as well as the proliferation of that technology. As high performance graphics workstations work cooperatively with supercomputers, and as real-time remote robotics and experimental control become a reality, the bandwidth requirements will continue to grow. In addition, as the number of sites on the networks increase, so will the aggregate bandwidth requirement. However, at the same time, the underlying bandwidth capabilities are also increasing. Satellite bandwidths of tens of megabits are available, and fiber optics technologies are providing extremely high bandwidths (in the range of gigabits). It is therefore essential that the underlying connectivity take advantage of these advances in communications to increase the available end-to-end bandwidth.

Expressway Routing

As higher levels of internet connectivity occur there will be a new set of problems related to lowest hop count and lowest delay routing metrics. The assumed internet connectivity can easily present situations where the highest speed, lowest delay route between two nodes on the same net is via a route on another network. Consider two sites one either end of the country, but both on the same multipoint internet, where their network also is gatewayed to some other network with high speed transcontinental links. The routing algorithms must be able to handle these situations gracefully, and they become of increased importance in handling global type-of-

service routing.

3. NETWORK SPECIFICATIONS

To achieve the end-to-end user functions discussed in section 2, it is not adequate to simply provide the underlying connectivity described in the previous section. The network must provide a certain set of capabilities on an end-to-end basis. In this section, we discuss the specifications on the network that are required.

Near Term Specifications

In the near term, the requirements on the networks are two-fold. First is to provide those functions that will permit full interoperability, and second the internetwork must address the additional requirements that arise in the connection of networks, users, and resources.

Interoperability

A first-order requirement for scientific computer networks (and computer networks in general) is that they be interoperable with each other, as discussed in the above section on connectivity. A first step to accomplish this is to use IP. The use of IP will allow individual networks built by differing agencies to combine resources and minimize cost by avoiding the needless duplication of network resources and their management. However, use of IP does not provide end-to-end interoperability. There must also be compatibility of higher level functions and protocols. At a minimum, while commonly agreed upon standards (such as the ISO developments) are proceeding, methods for interoperability between different protocol suites must be developed. This would provide interoperability of certain functions, such as file transfer, electronic mail and remote login. The emphasis, however, should be on developing agreement within the scientific community on use of a standard set of protocols.

Access Control

The design of the network should include adequate methods for controlling access to the network by unauthorized personnel. This especially includes access to network capabilities that are reachable via the commercial phone network and public data nets. For example, terminal servers that allow users to dial up via commercial phone lines should have adequate authentication mechanisms in place to prevent access by unauthorized individuals. However, it should be noted that most hosts that are reachable via such networks are also reachable via other "non-network" means, such as directly dialing

over commercial phone lines. The purpose of network access control is not to insure isolation of hosts from unauthorized users, and hosts should not expect the network itself to protect them from "hackers".

Privacy

The network should provide protection of data that traverses it in a way that is commensurate with the sensitivity of that data. It is judged that the scientific requirements for privacy of data traveling on networks does not warrant a large expenditure of resources in this area. However, nothing in the network design should preclude the use of link level or end-to-end encryption, or other such methods that can be added at a later time. An example of this kind of capability would be use of KG-84A link encryptors on MILNET or the Fig Leaf DES-based end-to-end encryption box developed by DARPA.

Accounting

The network should provide adequate accounting procedures to track the consumption of network resources. Accounting of network resources is also important for the management of the network, and particularly the management of interconnections with other networks. Proper use of the accounting database should allow network management personnel to determine the "flows" of data on the network, and the identification of bottlenecks in network resources. This capability also has secondary value in tracking down intrusions of the network, and to provide an audit trail if malicious abuse should occur. In addition, accounting of higher level network services (such as terminal serving) should be kept track of for the same reasons.

Type of Service Routing

Type of service routing is necessary since not all elements of network activity require the same resources, and the opportunities for minimizing use of costly network resources are large. For example, interactive traffic such as remote login requires low delay so the network will not be a bottleneck to the user attempting to do work. Yet the bandwidth of interactive traffic can be quite small compared to the requirements for file transfer and mail service which are not response time critical. Without type of service routing, network resources must sized according to the largest user, and have characteristics that are pleasing to the most finicky user. This has major cost implications for the network design, as high-delay links, such as satellite links, cannot be used for interactive traffic despite the significant cost savings they represent over terrestrial links. With type of service routing in place in the network gateways, and proper software in the hosts to make use of such

capabilities, overall network performance can be enhanced, and sizable cost savings realized. Since the IP protocol already has provisions for such routing, such changes to existing implementations does not require a major change in the underlying protocol implementations.

Administration of Address Space

Local administration of network address space is essential to provide for prompt addition of hosts to the network, and to minimize the load on backbone network administrators. Further, a distributed name to address translation service also has similar advantages. The DARPA Name Domain system currently in use on the Internet is a suitable implementation of such a name to address translation system.

Remote Procedure Call Libraries

In order to provide a standard library interface so that distributed network utilities can easily communicate with each other in a standard way, a standard Remote Procedure Call (RPC) library must be deployed. The computer industry has lead the research community in developing RPC implementations, and current implementations tend to be compatible within the same type of operating system, but not across operating systems. Nonetheless, a portable RPC implementation that can be standardized can provide a substantial boost in present capability to write operating system independent network utilities. If a new RPC mechanism is to be designed from scratch, then it must have enough capabilities to lure implementors away from current standards. Otherwise, modification of an existing standard that is close to the mark in capabilities seems to be in order, with the cooperation of vendors in the field to assure implementations will exist for all major operating systems in use on the network.

Remote Job Entry (RJE)

The capabilities of standard network RJE implementations are inadequate, and are implemented prolifically among major operating systems. While the notion of RJE evokes memories of dated technologies such as punch cards, the concept is still valid, and is favored as a means of interaction with supercomputers by science users. All major supercomputer manufacturers support RJE access in their operating systems, but many do not generalize well into the Internet domain. That is, a RJE standard that is designed for 2400 baud modem access from a card reader may not be easily modifiable for use on the Internet. Nonetheless, the capability for a network user to submit a job from a host and have its output delivered on a printer attached to a different host would be welcomed by most science users. Further, having this capability interoperate with

existing RJE packages would add a large amount of flexibility to the whole system.

Multiple Virtual Connections

The capability to have multiple network connections open from a user's workstation to remote network hosts is an invaluable tool that greatly increases user productivity. The network design should not place limits (procedural or otherwise) on this capability.

Network Operation and Management Tools

The present state of internet technology requires the use of personnel who are, in the vernacular of the trade, called network "wizards," for the proper operation and management of networks. These people are a scarce resource to begin with, and squandering them on day to day operational issues detracts from progress in the more developmental areas of networking. The cause of this problem is that a good part of the knowledge for operating and managing a network has never been written down in any sort of concise fashion, and the reason for that is because networks of this type in the past were primarily used as a research tool, not as an operational resource. While the usage of these networks has changed, the technology has not adjusted to the new reality that a wizard may not be nearby when a problem arises. To insure that the network can flexibly expand in the future, new tools must be developed that allow non-wizards to monitor network performance, determine trouble spots, and implement repairs or 'work-arounds'.

Future Goals

The networks of the future must be able to support transparent access to distributed resources of a variety of different kinds. These resources will include supercomputer facilities, remote observing facilities, distributed archives and databases, and other network services. Access to these resources is to be made widely available to scientists, other researchers, and support personnel located at remote sites over a variety of internetted connections. Different modes of access must be supported that are consonant with the sorts of resources that are being accessed, the data bandwidths required and the type of interaction demanded by the application.

Network protocol enhancements will be required to support this expansion in functionality; mere increases in bandwidth are not sufficient. The number of end nodes to be connected is in the hundreds of thousands, driven by increasing use of microprocessors and workstations throughout the community. Fundamentally different sorts of services from those now offered are anticipated, and dynamic

bandwidth selection and allocation will be required to support the different access modes. Large-scale internet connections among several agency size internets will require new approaches to routing and naming paradigms. All of this must be planned so as to facilitate transition to the ISO/OSI standards as these mature and robust implementations are placed in service and tuned for performance.

Several specific areas are identified as being of critical importance in support of future network requirements, listed in no particular order:

Standards and Interface Abstractions

As more and different services are made available on these various networks it will become increasingly important to identify interface standards and suitable application abstractions to support remote resource access. These abstractions may be applicable at several levels in the protocol hierarchy and can serve to enhance both applications functionality and portability. Examples are transport or connection layer abstractions that support applications independence from lower level network realizations or interface abstractions that provide a data description language that can handle a full range of abstract data type definitions. Applications or connection level abstractions can provide means of bridging across different protocol suites as well as helping with protocol transition.

OSI Transition and Enhancements

Further evolution of the OSI network protocols and realization of large-scale networks so that some of the real protocol and tuning issues can be dealt with must be anticipated. It is only when such networks have been created that these issues can be approached and resolved. Type-of-service and Expressway routing and related routing issues must be resolved before a real transition can be contemplated. Using the interface abstraction approach just described will allow definition now of applications that can transition as the lower layer networks are implemented. Applications gateways and relay functions will be a part of this transition strategy, along with dual mode gateways and protocol translation layers.

Processor Count Expansion

Increases in the numbers of nodes and host sites and the expected growth in use of micro-computers, super-micro

workstations, and other modest cost but high power computing solutions will drive the development of different network and interconnect strategies as well as the infrastructure for managing this increased name space. Hierarchical name management (as in domain based naming) and suitable transport layer realizations will be required to build networks that are robust and functional in the face of the anticipated expansions.

Dynamic Binding of Names to Addresses

Increased processor counts and increased usage of portable units, mobile units and lap-top micros will make dynamic management of the name/address space a must. Units must have fixed designations that can be re-bound to physical addresses as required or expedient.

4. USER SERVICES

The user services of the network are a key aspect of making the network directly useful to the scientist. Without the right user services, network users separate into artificial subclasses based on their degree of sophistication in acquiring skill in the use of the network. Flexible information dissemination equalizes the effectiveness of the network for different kinds of users.

Near Term Requirements

In the near term, the focus is on providing the services that allow users to take advantage of the functions that the interconnected network provides.

Directory services

Much of the information necessary in the use of the network is for directory purposes. The user needs to access resources available on the network, and needs to obtain a name or address.

White Pages

The network needs to provide mechanisms for looking up names and addresses of people and hosts on the network. Flexible searches should be possible on multiple aspects of the directory listing. Some of these services are normally transparent to the user/host name to address translation for example.

Yellow Pages

Other kinds of information lookup are based on cataloging and classification of information about resources on the networks.

Information Sharing Services

Bulletin Boards

The service of the electronic bulletin board is the one-to-many analog of the one-to-one service of electronic mail. A bulletin board provides a forum for discussion and interchange of information. Accessibility is network-wide depending on the definition of the particular bulletin board. Currently the SMTP and UUCP protocols are used in the transport of postings for many bulletin boards, but any similar electronic mail transport can be substituted without affecting the underlying concept. An effectively open-ended recipient list is specified as the recipient of a message, which then constitutes a bulletin board posting. A convention exists as to what transport protocols are utilized for a particular set of bulletin boards. The user agent used to access the Bulletin Board may vary from host to host. Some number of host resources on the network provide the service of progressively expanding the symbolic mail address of the Bulletin Board into its constituent parts, as well as relaying postings as a service to the network. Associated with this service is the maintenance of the lists used in distributing the postings. This maintenance includes responding to requests from Bulletin Board readers and host Bulletin Board managers, as well as drawing the appropriate conclusions from recurring automatically generated or error messages in response to distribution attempts.

Community Archiving

Much information can be shared over the network. At some point each particular information item reaches the stage where it is no longer appropriately kept online and accessible. When moving a file of information to offline storage, a network can provide its hosts a considerable economy if information of interest to several of them need only be stored offline once. Procedures then exist for querying and retrieving from the set of offline stored files.

Shared/distributed file system

It should be possible for a user on the network to look at a

broadly defined collection of information on the network as one useful whole. To this end, standards for accessing files remotely are necessary. These standards should include means for random access to remote files, similar to the generally employed on a single computer system.

Distributed Databases and Archives

As more scientific disciplines computerize their data archives and catalogs, mechanisms will have to be provided to support distributed access to these resources. Fundamentally new kins of collaborative research will become possible when such resources and access mechanisms are widely available.

Resource Sharing Services

In sharing the resources or services available on the network, certain ancillary services are needed depending on the resource.

Access Control

Identification and authorization is needed for individuals, hosts or subnetworks permitted to make use of a resource available via the network. There should be consistency of procedure for obtaining and utilizing permission for use of shared resources. The identification scheme used for access to the network should be available for use by resources as well. In some cases, this will serve as sufficient access control, and in other cases it will be a useful adjunct to resource-specific controls. The information on the current network location of the user should be available along with information on user identification to permit added flexibility for resources. For example, it should be possible to verify that an access attempt is coming from within a state. A state agency might then grant public access to its services only for users within the state. Attributes of individuals should be codifiable within the access control database, for example membership in a given professional society.

Privacy

Users of a resource have a right to expect that they have control over the release of the information they generate. Resources should allow classifying information according to degree of access, i.e. none, access to read, access according to criteria specified in the data itself, ability to change or add information. The full range of identification information described under access control should be available to the user when specifying access. Access could be granted to all fellow members of a professional society, for example.

Accounting

To permit auditing of usage, accounting information should be provided for those resources for which it is deemed necessary. This would include identity of the user of the resource and the corresponding volume of resource components.

Legalities of Interagency Research Internet

To make the multiply-sponsored internetwork feasible, the federal budget will have to recognize that some usage outside a particular budget category may occur. This will permit the cross-utilization of agency funded resources. For example, NSFnet researchers would be able to access supercomputers over NASnet. In return for this, the total cost to the government will be significantly reduced because of the benefits of sharing network and other resources, rather than duplicating them.

Standards

In order for the networking needs of scientific computing to be met, new standards are going to evolve. It is important that they be tested under actual use conditions, and that feedback be used to refine them. Since the standards for scientific communication and networking are to be experimented with, they are more dynamic than those in other electronic communication fields. It is critical that the resources of the network be expended to promulgate experimental standards and maximize the range of the community utilizing them. To this end, the sharing of results of the testing is important.

User-oriented Documentation

The functionality of the network should be available widely without the costly need to refer requests to experts for formulation. A basic information facility in the network should therefore be developed. The network should be self-documenting via online help files, interactive tutorials, and good design. In addition, concise, well-indexed and complete printed documentation should be available.

Future Goals

The goal for the future should be to provide the advanced user services that allow full advantage to be taken of the interconnection of users, computing resources, data bases, and experimental facilities. One major goal would be the creation of a national knowledge bank. Such a knowledge bank would capture and organize computer-based knowledge in various scientific fields that is currently available only in written/printed form, or in the minds of

experts or experienced workers in the field. This knowledge would be stored in knowledge banks which will be accessible over the network to individual researchers and their programs. The result will be a codification of scientific understanding and technical know-how in a series of knowledge based systems which would become increasingly capable over time.

CONCLUSION

In this paper, we have tried to describe the functions required of the interconnected national network to support scientific research. These functions range from basic connectivity through to the provision for powerful distributed user services.

Many of the goals described in this paper are achievable with current technology. They require coordination of the various networking activities, agreement to share costs and technologies, and agreement to use common protocols and standards in the provision of those functions. Other goals require further research, where the coordination of the efforts and sharing of results will be key to making those results available to the scientific user.

For these reasons, we welcome the initiative represented by this workshop to have the government agencies join forces in providing the best network facilities possible in support of scientific research.

APPENDIX

Internet Task Force on Scientific Computing

Rick Adrion	University of Massachusetts
Ron Bailey	NASA Ames Research Center
Rick Bogart	Stanford University
Bob Brown	RIACS
Dave Farber	University of Delaware
Alan Katz	USC Information Science Institute
Jim Leighton	Lawrence Livermore Laboratories
Keith Lantz	Stanford University
Barry Leiner	(chair) RIACS
Milo Medin	NASA Ames Research Center
Mike Muuss	US Army Ballistics Research Laboratory
Harvey Newman	California Institute of Technology
David Roode	Intellicorp
Ari Ollikainen	General Electric
Peter Shames	Space Telescope Science Institute
Phil Scherrer	Stanford University